eJNDT e-Journal of Nondestructive Testing - ISSN 1435-4934 - www.ndt.net

EWSHM 2024
11th European Workshop on Structural Health Monitoring

Damage-aware structural control based on reinforcement learning

Aleksandra JEDLIŃSKA, Dominik PISARSKI, Grzegorz MIKUŁOWSKI, Bartłomiej BŁACHOWSKI, Łukasz JANKOWSKI

Institute of Fundamental Technological Research, Polish Academy of Sciences
Pawińskiego 5B, 02-106 Warsaw, Poland
e-mail: ljank@ippt.pan.pl

Abstract. This contribution presents a semi-active control technique intended for mitigation of structural vibrations. The control law is derived in a repeated trial-and-error interaction between the control agent and a simulated environment. The experience-based training approach is used which is the defining feature of the machine learning techniques of reinforcement learning (RL), implemented here using the framework provided by Deep Q Learning (DQN). The involved artificial neural network not only determines the control action, but additionally identifies structural damages, which is a nontrivial task due to the nonlinearity of the control. This requires a specific multi-head architecture, which allows the network to be damage-aware, and a specific training procedure, where the memory pool preserved for the RL stage of experience replay is populated with not only the observations, control actions, and rewards, but also with the momentary status of structural damage. Such an approach can be used to explicitly promote the damage-awareness of the control agent. The proposed technique is tested and verified in a numerical example of a shear-type building model subjected to a random seismic-type excitation. A tuned mass damper (TMD) with a controllable level of viscous damping is used to implement the semi-active actuation, and the optimally tuned classical TMD provides the reference response.

Keywords: semi-active control, tuned mass damper (TMD), reinforcement learning, damage identification

Introduction

Structural vibrations arise due to operational conditions or external excitations and can have very negative consequences for structural integrity. In the case of extreme seismic excitations, the resulting response is also extreme and can lead to structural collapse, large costs, and loss of life. Therefore, various techniques have been investigated, developed, and applied for the purpose of structural control and mitigation of structural vibrations [1]. As discussed in the following, they generally encompass a full spectrum of approaches that range

NDT.net Media and Publishing Partner
https://doi.org/10.58286/29606

from fully passive methods, through semi-active control approaches, to fully active control techniques.

Active control techniques are well-researched and very effective [2]. However, they rely on dedicated actuators that generate large external control forces. As a result, active control approaches require large power supply, which can be missing during a major earthquake, and in case of malfunction can potentially lead to structural instabilities.

Semi-active control techniques provide control that is often called dissipative control [3–5]. Instead of directly generating large external forces, these techniques rely on local modification of mechanical properties of selected structural members or actuators (like viscous damping, moment-bearing ability, etc.) [6]. Consequently, the employed actuators generate only dissipative forces, that is forces that oppose the movement of the structure. No large external power sources are thus needed, and the danger of instabilities is reduced. However, the ensuing control is nonlinear, and the resulting optimal control problems are often not solvable using analytical means.

Passive techniques for mitigation of vibrations rely usually on structural optimization and dedicated dissipative devices. In case of seismic vibrations, the most popular approaches [7] include base isolation [8] and tuned mass dampers [9]. The former aim to partially decouple the superstructure from the ground motion. The latter, the tuned mass damper, is a classical engineering device that dates back a century ago to the works of Frahm [10], as well as Ormondroyd and Den Hartog [11]. In various implementations, it has been physically installed and proved effective in numerous high-rise buildings around the world [12].

This study considers a tuned mass damper (TMD) as an actuator, and its application and effectiveness in case of a damaged superstructure. In theoretical terms, a TMD consists of a mass connected to the structure (in case of high-rise buildings, usually at an upper story), through a spring and a dashpot. Several modifications to the classical TMD concept have been proposed and investigated [13]. They often turn a passive TMD into an actively or semi-actively controlled device. Typical active approaches involve an additional actuator placed in-between the mass and the structure that can generate additional control forces [14]. This contribution focuses on semi-actively controlled TMDs. The considered control relies on real-time tuning of the dashpot damping [15]. Formally, such a control is bilinear, and analytical solutions for the optimal control are rarely available. However, classical results in optimal control theory [16] show that the open loop optimal control is often of the bang-bang type, that is, it switches between two extreme states. This assumption is adopted here, and the switching points are learned using the machine learning approach of reinforcement learning (RL) [17-19], that is, through a series of trial-and-error interactions with a simulated structure subjected to a random seismic-type excitation. The control obtained this way is assessed in terms of the root mean square displacement of the top floor. It is demonstrated to be significantly more effective than that of the optimally tuned passive TMD.

The approach proposed here builds upon our earlier results presented in [20,21]. However, the case of a damaged structure is considered here. The RL control agent is trained in a data-driven manner using a model of the undamaged structure. This work deals with the three questions that naturally emerge when such a control is applied to a damaged structure:

1. Can the damage be identified based on the measured responses of the RL-controlled structure?
2. Does the RL control remain effective when the structure is damaged?
3. Can the knowledge about the identified damage be used by the RL control agent to increase the effectiveness of the control?

Section 1 introduces the modelled structure. Section 2 presents the effectiveness of the RL control and compares it to that of the reference optimally tuned passive TMD. Section 3 considers the problem of damage identification in such a data-driven control system. Finally, Section 4 proposes a damage-aware RL control agent and tests it effectiveness.

1. The Structure

1.1 Shear Type Building

Reinforcement learning (RL) requires trial-and-error type interaction between the control agent being trained and the structure. However, this kind of interaction cannot be implemented with real-world physical structures due to safety reasons. The training phase must be thus implemented in interaction with a numerically simulated structure. Dynamics of such a structure is only an approximation to the dynamics of the modelled physical structure. A careful consideration of the robustness of the derived control with respect to model errors is therefore necessary. As demonstrated in [21], the control approach developed here is relatively insensitive to these errors.

The structure considered here modelled an 11-story shear building, with a single degree of freedom (DOF) per story, equipped with a TMD at the top. A schematic depiction is shown in Figure 1.

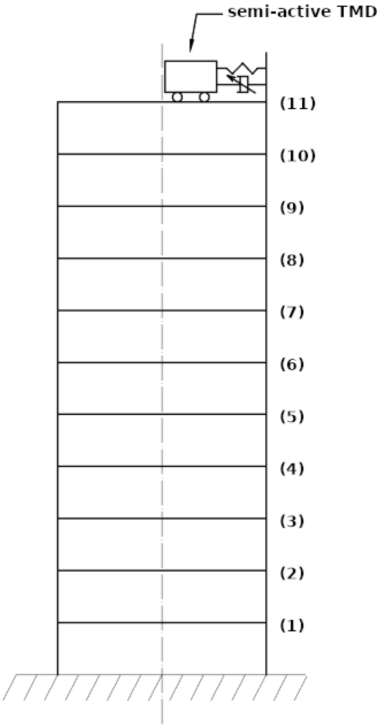

Fig. 1. A schematic diagram of the 11-story shear building considered in this study. The semi-active TMD is mounted at the top story.

A seismic type excitation was considered and modelled as a lateral acceleration $a(t)$. The equation of motion is thus formulated in 12 DOFs. It takes the following form:

$$\mathbf{M}\ddot{\mathbf{x}}(t) + \mathbf{C}(t)\dot{\mathbf{x}}(t) + \mathbf{K}\mathbf{x}(t) = -\mathbf{M}\mathbf{r}a(t)$$

where the mass and stiffness matrices are denoted by $\mathbf{M}$ and $\mathbf{K}$, and $\mathbf{C}(t)$ is the time-dependent damping matrix. The considered bang-bang control amounts to switching $\mathbf{C}(t)$ between two extreme states that, besides the standard damping of the building, correspond to the maximum and zero level of TMD dashpot damping. The lateral load allocation vector $\mathbf{r}$ is a 12-element vector of ones. For the building without the TMD, a material damping model was assumed with 2% critical damping ratio at the first natural vibration mode. The masses and stiffnesses of each story follow an example structure already considered in literature [22]. The mass of the TMD is 3% of the building mass. The natural frequencies range from 0.887 Hz to 14.917 Hz. For the purpose of response simulation, the Newmark algorithm was used with the time step of 2 ms. The maximum control switching frequency was 25 Hz.

1.2 Seismic Acceleration and Feedback Measurements

RL training typically proceeds in episodes. To avoid overfitting the control agent to a specific ground motion pattern, the seismic excitation was modelled as a white noise anew in each episode. In all time steps, the momentary ground acceleration value $a(t)$ was independently drawn from a uniform probability distribution centred around zero.

The feedback signal consisted of the full state vector. Such a choice is inconvenient in practice, but it corresponds to the full observation and allows thus to assess the ultimate potential of the considered control and RL agent architecture.

2. Reinforcement Learning (RL) Control for Undamaged Building

The RL control was implemented using the double DQN framework [23]. The control objective was expressed using the reward, defined as the total instantaneous mechanical energy of the building structure (without the TMD) and processed using the smooth L1 loss.

2.1 RL Agent Architecture

Figure 2 presents the architecture of the artificial neural network used by the RL agent to select control actions (TMD damping on/off). The input layer corresponds to the feedback signal (state vector), while the output layer corresponds to the two possible control actions. Network hyperparameters (number/type of layers, number of neurons, etc.) were selected using the trial-and-error approach, and, in order to safeguard some spare network capacity, they were kept slightly above the minimum level at which the efficiency started to deteriorate. AdamW optimizer was used with the initial learning rate of 0.001. The training episodes were 20 s long, which corresponded to 18 fundamental periods of vibration. The initial exploration rate was 100% and decreased to 5% with the factor of 0.994 per episode.

input(24)
2x (dense(60) + Leaky ReLU)
2x (dense(60) + Leaky ReLU)
dense(2)

Fig. 2. Architecture of the artificial neural network in the RL control agent.

2.2 Control Effectiveness

The effectiveness of the developed control was quantified in the same terms as the reward. The mean total mechanical energy of the building structure (not counting the TMD, averaged over the entire test episode) was thus used and compared to the reference provided by the classical, optimally tuned passive TMD. For evaluation, 4000 test episodes were simulated, each 80 s long. The results are listed in Table 1. The RL control was significantly more effective than the optimally tuned passive TMD: it resulted in the mean total energy of the building being reduced by 25% on average. Additionally, the RL control reduced the standard deviation of the mean total energy to an even greater extent, that is, by 41.7%.

Table 1. Effectiveness of RL control agent trained using the undamaged structure (Fig. 2), compared to the optimally tuned passive TMD, in application to the undamaged structure. Mean total mechanical energy of the building (without the TMD) and energy reduction per single episode.

	mean	std. dev.
with optimum passive TMD [a.u.]	667.089	168.694
with RL control [a.u.]	491.022	98.337
energy reduction per episode (TMD→RL)	24.977 %	9.128 %

3. Damage Identification for RL-controlled Building

As demonstrated in the previous section, the RL control is significantly more effective than the optimally tuned passive TMD. However, the control provided by the trained RL agent is data-driven and nonlinear, and consequently, the dynamics of the RL- controlled building is effectively also nonlinear. This hinders application of typical structural health monitoring methods, which are based on the assumption of linearity of structural response. A machine learning approach is thus used here for damage identification in the investigated RL-controlled structure.

Structural damage was modelled as stiffness reduction (up to 50%) at up to four randomly selected building floors (but not the TMD). At each selected floor, the stiffness reduction ratio was randomly and independently drawn from the uniform distribution $U(0\%, 50\%)$. In this way, the damage scenario was generated anew for each training and testing episode.

For the purpose of damage identification, an additional branch was created in the neural network architecture of the RL-agent, as shown in Figure 3. The 11 elements of its output layer correspond to the 11 floors of the building, and they are intended to estimate the respective damage reduction ratios. The input layer is the feedback signal, and it is shared with the RL-control branch.

input(24)		
2x (dense(60) + Leaky ReLU)		2x (lstm(60) + tanh)
2x (dense(60) + Leaky ReLU)		2x (dense(60) + Leaky ReLU)
dense(2)		dense(11)

Fig. 3. Architecture of the artificial neural network in the RL control agent with damage identification branch.

During training, the control branch was used for deciding the control action, and its weights were frozen. The damage identification branch was trained in interaction with the damaged structure, using the mean square identification error as the loss. The trained network was assessed using 4000 test episodes, each 80 s long, and each with its own damage scenario. Figure 4 presents the rms errors of the identified damage reduction ratios (across all time steps and all test episodes). The error magnitudes were similarly small across all floors and ranged between 1.2 and 2.4%. These values should be related to the magnitudes of the actually simulated damages, which were uniformly distributed between 0 and 50% for each damaged floor. In comparison, results generated for each story uniformly at random would yield the rms of 26.5%, while the rms of the best constant deterministic prediction would be 14.2%.

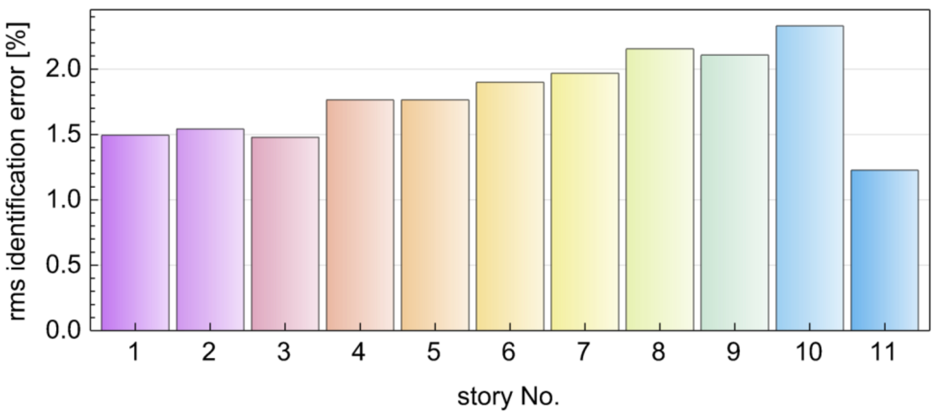

Fig. 4. Root mean square errors of damage identification ratios for individual floors.

4. Damage-aware RL Control

4.1 Deterioration of Control Effectiveness in the Damaged Structure

While training the damage identification branch of the RL agent, the control branch was frozen. The same control law was thus applied with the damaged building as with the original undamaged building, even though their dynamic characteristics are different. Expectably, this negatively affected the control effectiveness, as seen in Table 2 (which partially repeats the data from Table 1). For the TMD- and RL-controlled structures, the mean total energy per episode has increased on average by 10.1% and 7.45%, respectively.

Table 2. Effectiveness of RL control agent trained using undamaged structure (Fig. 2 and 3), compared to the optimally tuned passive TMD, in application to undamaged and damaged structures. Mean total mechanical energy of the building (without the TMD) and energy reduction per single episode.

	undamaged structure		damaged structure	
	mean	std. dev.	mean	std. dev.
with optimum passive TMD [a.u.]	667.089	168.694	734.261	220.322
with RL control [a.u.]	491.022	98.337	527.585	118.318
energy reduction per episode (TMD→RL)	24.977 %	9.128 %	26.046 %	10.863 %

4.2 Damage-Aware Control

Two approaches have been tested to increase the control effectiveness:

1. The entire control branch was fine-tuned in interaction with damaged structures. Each training episode used a randomly generated damage scenario.
2. Before fine-tuning the control branch in interaction with damaged structures, its head (last 3 layers) was additionally fed with damage-related features extracted in the damage identification branch, as shown in Figure 4. This allowed the control head to be damage-aware, that is, to take into account the damage-related information while processing the control-related features.

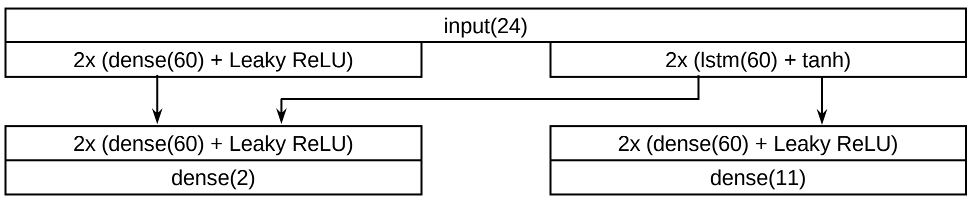

Fig. 4. Architecture of the artificial neural network in the RL control agent with damage identification branch and damage-related features fed into the head of the control branch.

Only the second approach led to a noticeable improvement in control effectiveness. While fine-tuning the control branch, all other network weights were frozen. The resulting network was assessed using 4000 test episodes, each 80 s long, and each with its own independently generated random damage scenario. The results are listed in Table 3, and they should be compared to these in Table 2 above. A graphical representation of all cases is presented in Figure 5.

Table 3. Effectiveness of RL control agent fine-tuned using damaged structures (Fig. 4), compared to the optimally tuned passive TMD, in application to damaged structures. Mean total mechanical energy of the building (without the TMD) and energy reduction per single episode.

	mean	std. dev.
with optimum passive TMD [a.u.]	732.412	222.184
with RL control [a.u.]	503.329	118.531
energy reduction per episode (TMD→RL)	28.429 %	8.445 %

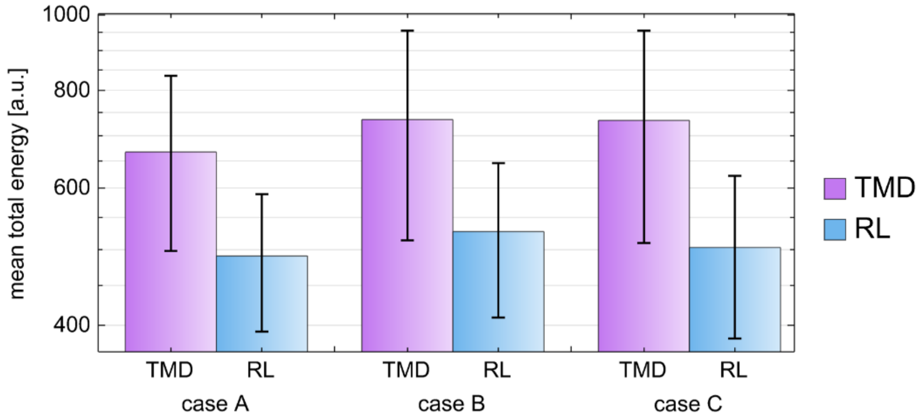

Fig. 5. Mean total energy of the structure (without the TMD) for the RL control, compared to that of the optimally tuned passive TMD. Case A: RL agent trained and tests performed on the undamaged structure. Case B: RL agent trained on the undamaged structure; tests performed using damaged structures. Case C: RL agent additionally fine-tuned and tests performed using damaged structures. Whiskers denote the $\pm 1\sigma$ range.

Application of the original RL control to the damaged structure resulted in the undesirable increase of the mean total energy by 7.45%. Fine-tuning the control branch with the damage-related features reduced this increase to 2.51% only.

Conclusion

This contribution presented and evaluated a multi-head artificial neural network for a DQN reinforcement learning agent designed for the double task of semi-active control a shear-type building model (subjected to random seismic type excitation) and identification of its damages (modelled at up to 4 floors as stiffness reductions of up to 50% each). The control branch within the network, besides the control-related features, processes also the damage-related features obtained in the damage identification branch. Such an architecture allows the control task to be aware of damages and to accordingly modify its control actions. The main findings can be summarized as follows:

1. In comparison to the same building model equipped with an optimally tuned passive TMD, the RL agent provided control that reduced the mean total mechanical energy by 25% on average. Additionally, the standard deviation of the total energy was also significantly reduced (by 42%).
2. In tests with damaged structures, the designed RL agent was able to identify damage factors at all 11 floors with the rms error of 1.8% of the original story stiffness. This is a much better accuracy than a uniform random guess (26.5%) and a constant deterministic guess (14.2%).
3. Re-using the damage-related features in the control branch significantly improved the control. When the original RL control was applied to the damaged structure, the mean total energy increased by 7.45%. Re-using damage features for control purposes reduced this increase to 2.51% only.

Acknowledgement

The authors gratefully acknowledge the support of the National Science Centre, Poland, granted under the grant agreement 2020/39/B/ST8/02615. For the purpose of Open Access, the authors have applied a CC-BY public copyright licence to any Author Accepted Manuscript (AAM) version arising from this submission.

References

[1] Basu B., Bursi O.S., Cascati F., Cascati S., Del Grosso A.E., Domaneshi M., Faravelli L., Holnicki-Szulc J., Irshik H., Krommer M., Lepidi M., Martelli A., Ozturk B., Pozo F., Pujol G., Rakicevic Z., Rodellar J. A European Association for the Control of Structures joint perspective. Recent studies in civil structural control across Europe (2014) Structural Control and Health Monitoring, 21:1414–1436.

[2] Korkmaz S. A review of active structural control: Challenges for engineering informatics (2011) Computers and Structures, 89 (23–24): 2113 – 2132. https://doi.org/10.1016/j.compstruc.2011.07.010

[3] Fisco N.R., Adeli H. Smart structures: Part I – Active and semi-active control (2011) Scientia Iranica, 18 (3A): 275–284. https://doi.org/10.1016/j.scient.2011.05.034

[4] Casciati F., Rodellar J., Yildirim U. Active and semi-active control of structures-theory and applications: A review of recent advances (2012) Journal of Intelligent Material Systems and Structures, 23 (11):1181–1195. https://doi.org/10.1177/1045389X12445029

[5] Symans M.D., Constantinou M.C. Semi-active control systems for seismic protection of structures: A state-of-the-art review (1999) Engineering Structures, 21 (6):469–487. https://doi.org/10.1016/S0141-0296(97)00225-3

[6] Hurlebaus S., Gaul L. Smart structure dynamics (2006) Mechanical Systems and Signal Processing, 20 (2):255–281. https://doi.org/10.1016/j.ymssp.2005.08.025

[7] Parulekar Y.M., Reddy G.R. Passive response control systems for seismic response reduction: A state-of-the-art review (2009) International Journal of Structural Stability and Dynamics, 9 (1):151–177. https://doi.org/ 10.1142/S0219455409002965

[8] Warn G.P., Ryan K.L. A review of seismic isolation for buildings: Historical development and research needs (2012) Buildings, 2 (3):300–325. https://doi.org/10.3390/buildings2030300

[9] Gutierrez Soto M., Adeli H. Tuned Mass Dampers (2013) Archives of Computational Methods in Engineering, 20 (4):419–431. https://doi.org/ 10.1007/s11831-013-9091-7

[10] Frahm H. Device for damping vibrations of bodies (1911) U.S. Patent No. 989,958.

[11] Ormondroyd J., Den Hartog J.P. The Theory of the Dynamic Vibration Absorber (1928) Journal of Fluids Engineering, Transactions of the ASME, 49–50 (2):021007. https://doi.org/10.1115/1.4058553

[12] Elias S., Matsagar V. Research developments in vibration control of structures using passive tuned mass dampers (2017) Annual Reviews in Control, 44:129–156. https://doi.org/10.1016/j.arcontrol.2017.09.015

[13] Nagarajaiah S., Jung H.-J. Smart tuned mass dampers: Recent developments (2014) Smart Structures and Systems, 13 (2):173–176. https://doi.org/10.12989/sss.2014.13.2.173

[14] Nishimura I., Kobori T., Sakamoto M., Koshika N., Sasaki K., Ohrui S. Active tuned mass damper (1992) Smart Materials and Structures, 1 (4):306–311. https://doi.org/10.1088/0964-1726/1/4/005

[15] Pinkaew T., Fujino Y. Effectiveness of semi-active tuned mass dampers under harmonic excitation (2001) Engineering Structures, 23 (7):850–856. https://doi.org/10.1016/S0141-0296(00)00091-2

[16] Kirk, D.E. Optimal control theory: an introduction (2004) Courier Corporation.

[17] Silver D., Hubert T., Schrittwieser J., Antonoglou I., Lai M., Guez A., Lanctot M., Sifre L., Kumaran D., Graepel T., Lillicrap T., Simonyan K., Hassabis D. A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play (2018) Science 362(6419):1140–1144, https://doi.org/10.1126/science.aar6404

[18] Khalatbarisoltani A., Soleymani M., Khodadadi M. Online control of an active seismic system via reinforcement learning (2019) Structural Control and Health Monitoring 26(3):e2298. https://doi.org/10.1002/stc.2298

[19] Pisarski D., Jankowski Ł. Reinforcement learning-based control to suppress the transient vibration of semi-active structures subjected to unknown harmonic excitation (2023) Computer-Aided Civil and Infrastructure Engineering 38 (12):1605–1621. http://doi.org/10.1111/mice.12920

[20] Jedlińska A., Pisarski D., Mikułowski G., Błachowski B., Jankowski Ł., Semi-active structural control using viscous dampers and reinforcement learning (2023) Proceedings of the 10th ECCOMAS Thematic Conference on Smart Structures and Materials (SMART 2023), July 3–5 2023, Patras (Greece), pp. 589–596. https://doi.org/10.7712/150123.9813.444767

[21] Jedlinska A., Pisarski D., Mikulowski G., Blachowski B., Jankowski Ł. Semi-Active Control of a Shear Building Based on Reinforcement Learning: Robustness to Measurement Noise and Model Error (2023) Proceedings of the 18th Conference on Computer Science and Intelligence Systems (FedCSIS 2023), Sept 17–20 2023, Warsaw (Poland), pp. 1007–1010. https://doi.org/10.15439/2023F8946

[22] Pourzeynali S., Lavasani H.H., Modarayi A.H., Active control of high rise building structures using fuzzy logic and genetic algorithms (2007) Engineering Structures 29(3):346–357. https://doi.org/ 10.1016/j.engstruct.2006.04.015

[23] Van Hasselt H., Guez A., Silver D. Deep reinforcement learning with double Q-Learning (2016) Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI 2016), Feb 12–17 2016, Phoenix (AZ, US), pp. 2094–2100. https://doi.org/ 10.1609/aaai.v30i1.10295